arXiv:1409.4324v2 [stat.ME] 7 Nov 2014

Improving Inference of Gaussian Mixtures Using Auxiliary Variables

Andrea Mercatanti [1] Fan Li [2] Fabrizia Mealli [3]

ABSTRACT

Expanding a lower-dimensional problem to a higher-dimensional space and then projecting back is often beneficial. This article rigorously investigates this perspective in the context of finite mixture models, namely how to improve inference for mixture models by using auxiliary variables. Despite the large literature in mixture models and several empirical examples, there is no previous work that gives general theoretical justification for including auxiliary variables in mixture models, even for special cases. We provide a theoretical basis for comparing inference for mixture multivariate models with the corresponding inference for marginal univariate mixture models. Analytical results for several special cases are established. We show that the probability of correctly allocating mixture memberships and the information number for the means of the primary outcome in a bivariate model with two Gaussian mixtures are generally larger than those in each univariate model. Simulations under a range of scenarios, including misspecified models, are conducted to examine the improvement. The method is illustrated by two real applications in ecology and causal inference.

KEY WORDS: bivariate, EM, Gaussian, information matrix, mixture model, score function

[1]Andrea Mercatanti is researcher, Statistics Department, Bank of Italy, Rome, Italy. (email: mercatan@libero.it).

[2]Fan Li is assistant professor, Department of Statistical Science, Duke University, Durham, NC, USA. (email: fli@stat.duke.edu).

[3]Fabrizia Mealli is professor, Department of Statistics, Informatics, Applications, University of Florence, Italy. (email: mealli@ds.unifi.it).

1 Introduction

The idea of expanding a lower-dimensional problem to a higher-dimensional space and then projecting back has been used in statistics and other disciplines. This article discusses a specific example in the context of finite mixture models; in particular, we rigorously investigate the impact on inference for mixture models when using auxiliary variables. Finite mixture models are a large class of statistical models for studying a wide variety of practical problems; comprehensive reviews can be found in McLachlan and Basford (1988); McLachlan and Peel (2000). The common idea underlying these models is that data are obtained from two or more underlying populations with common distributional form but different parameters. Formally, the data $\mathbf{x} = (\mathbf{x}_1, ..., \mathbf{x}_n)'$, where $\mathbf{x}_i$ is a m-dimensional vector and n is the sample size, follow the distribution:

$$f(\mathbf{x}) = \sum_{k=1}^{K} p_k f_k(\mathbf{x}; \boldsymbol{\theta}_k), \tag{1}$$

with the weights p_k's satisfying $p_k > 0$ and $\sum_k p_k = 1$. Standard choices of the densities f_k include Gaussian (normal), Poisson and Student's t-distributions. Here we focus on the most widely used Gaussian mixture models.

A main source of uncertainty in estimating mixture models is attributed to the unknown mixture membership of each unit. The EM algorithm (Dempster et al., 1977), which augments the mixture membership for each unit iteratively, is the most common approach for deriving maximum likelihood estimates (MLEs) of the parameters in mixture models. For estimating the variance matrix of the ML estimator, there are three main approaches. The first involves the "complete-data" likelihood, where the augmented mixture membership for each unit is treated as observed (e.g. Louis, 1982). The second are resampling-based methods (e.g. Newton and Raftery, 1994; Basford et al., 1997). The third is based on the original "incomplete-data" likelihood (e.g. Dietz and Böhning, 1996). An important recent work in this area is Boldea

and Magnus (2009), who derived the analytical forms of the score vector and Hessian matrix for Gaussian mixture models with arbitrary number of components and dimension of observations. Besides the likelihood-based approaches, there is also a large literature on the Bayesian approach to mixture models (e.g. West, 1992; West et al., 1994; Richardson and Green, 1997; Marin et al., 2005, and references therein).

Regardless of the mode of inference used, the key to inference for mixture models is to disentangle the unknown mixtures. Our main message here is that inference for mixture models can be sharpened by jointly modeling the primary variable with available auxiliary variable(s). Despite the large literature on multivariate mixture models, cross-dimensional comparison is rare. Multivariate analysis is usually conducted when the features of several variables or the relationship between the variables is of interest, but it is seldom considered for the purpose of sharpening univariate inference. Indeed, it is not obvious why including auxiliary variables in the models would improve estimation of the parameters for the primary variable. Clearly jointly modeling the primary variable with any arbitrary random variable would in principle only increase noise. But in real applications, (auxiliary) variables are usually associated with the mixture membership. Thus, on one hand, proper utilization of those relevant auxiliary variables may provide extra information to predict the mixture membership and consequently to disentangle the mixtures. On the other hand, however, for a given sample size modeling auxiliary variables could induce extra uncertainty because of the estimation of additional parameters; further, it increases model complexity and thus the risk of mis-specification. We show that the potential benefits dominate the potential drawbacks.

There are a few empirical examples within specific contexts that display the benefit of using auxiliary variables in mixture models. In the context of causal inference, a common goal in randomized clinical trials is to evaluate the effect of a drug or a therapy on a primary clinical outcome. While measurements on other features, such as side effects, are routinely collected,

they are usually analyzed separately, one at a time. When noncompliance arises, mixture models are often used since the population is heterogenous regarding compliance behavior (e.g. Imbens and Rubin, 1997). Mattei et al. (2013) and Mealli and Pacini (2013) show that jointly modeling primary and secondary outcomes significantly sharpens the inference for the primary outcome. Another example is found in the context of small-area estimation: DeSouza (1992) showed that analysis based on bivariate hierarchical models, which can be viewed as a special case of mixture models, reduces the posterior standard errors of the mean small-area estimates compared to those based on univariate models. However, to our knowledge, there is no previous work that gives general theoretical justification for this practice or explains the reasons underlying the benefit of auxiliary variables in mixture models, even for special cases.

The goal of this article is to provide a theoretical basis for comparing inference for multivariate mixture models with inference for the corresponding marginal univariate mixture models, filling a gap in the literature. In particular, proceeding from the incomplete-data likelihood perspective, we will establish analytical results for several special cases, showing that multivariate analysis increases the probability of correctly allocating the mixture membership and improves precision (or equivalently, reduces standard errors) of the maximum likelihood (ML) estimates compared to the corresponding univariate analysis. Another key insight from our results, partly shown in our empirical analysis, is that the introduction of an auxiliary variable tends to regularize the model and thus reduce the prevalence and the likelihood of spurious roots. We show these benefits clearly dominate the extra uncertainty due the larger parameters set involved by the auxiliary variable. As closed-form arguments on general mixture models are difficult to obtain, our analytical derivations are focused on the simple case of bivariate mixture models with two Gaussian components; models with mis-specification (non-Gaussian) and higher dimensions are explored in simulations and real applications.

The rest of the paper is organized as follows. In Section 2, we illustrate the intuition by

a simple visual example and present the main theoretical results. In Section 3, we conduct simulations to examine the small-sample comparisons between bivariate and univariate analyses under a variety of settings. Two real applications are presented in Section 4. Section 5 concludes.

2 Comparing univariate and bivariate mixture models

2.1 Basic setup and intuition

Consider a mixture model of two Gaussian densities,

$$f(\mathbf{x}) = p\, f_1(\mathbf{x}) + (1-p)\, f_2(\mathbf{x}), \tag{2}$$

where $f_k(\mathbf{x}) = |\boldsymbol{V}_k|^{-1/2} \exp\left\{-(\mathbf{x}-\boldsymbol{\mu}_k)'\boldsymbol{V}_k^{-1}(\mathbf{x}-\boldsymbol{\mu}_k)/2\right\}/2\pi$ for $k = 1, 2$. For a univariate density, $\mathbf{x} = x_1, \boldsymbol{\mu}_k = \mu_{1k}, \boldsymbol{V}_k = \sigma_{1k}^2$, while for a bivariate density,

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix},\ \boldsymbol{\mu}_k = \begin{pmatrix} \mu_{1k} \\ \mu_{2k} \end{pmatrix},\ \boldsymbol{V}_k = \begin{pmatrix} \sigma_{1k}^2 & \rho_k\sigma_{1k}\sigma_{2k} \\ \rho_k\sigma_{1k}\sigma_{2k} & \sigma_{2k}^2 \end{pmatrix}. \tag{3}$$

In what follows, we will use the subscript $m(=1,2)$ to denote the outcome and $k(=1,2)$ to denote the mixture component.

The intuition of the benefit of using the second outcome can be illustrated by a simple example in Figure 1. Consider four sets of parameters in (3), all with $p = 0.5$, $\boldsymbol{V}_1 = \boldsymbol{V}_2$ and $\sigma_1 = \sigma_2 = 1$, but different means and correlations: (a) $\boldsymbol{\mu}_1 = (0,0)', \boldsymbol{\mu}_2 = (0.05,1)', \rho = 0$; (b) $\boldsymbol{\mu}_1 = (0,0)', \boldsymbol{\mu}_2 = (0.05,1)', \rho = 0.9$; (c) $\boldsymbol{\mu}_1 = (0,0)', \boldsymbol{\mu}_2 = (0.05,4)', \rho = 0$; (d) $\boldsymbol{\mu}_1 = (0,0)', \boldsymbol{\mu}_2 = (0.05,4)', \rho = 0.9$. Figure 1 displays the empirical contour plots from 1000 samples generated from the above four settings. In all the settings, the underlying marginal distribution of x_1 is the same, very close to a standard Gaussian. Thus it would be difficult to disentangle the components based on a univariate analysis on x_1 alone. In contrast, in the

presence of a small distance between the means of x_2 in the two components, as in settings (a) and (b), there is already a mild but noticeable improvement in the separation of the mixtures, reflected by the bend in the contour lines near the middle in Figures 1 (a) and (b). When the distance increases, as in settings (c) and (d), the separation of the components becomes very visible. This is most striking in Figure 1(d), where the two components are completely separated. Given the same distance between the means of x_2 in the two components, higher conditional correlation within each component also appears to improve the disentanglement.

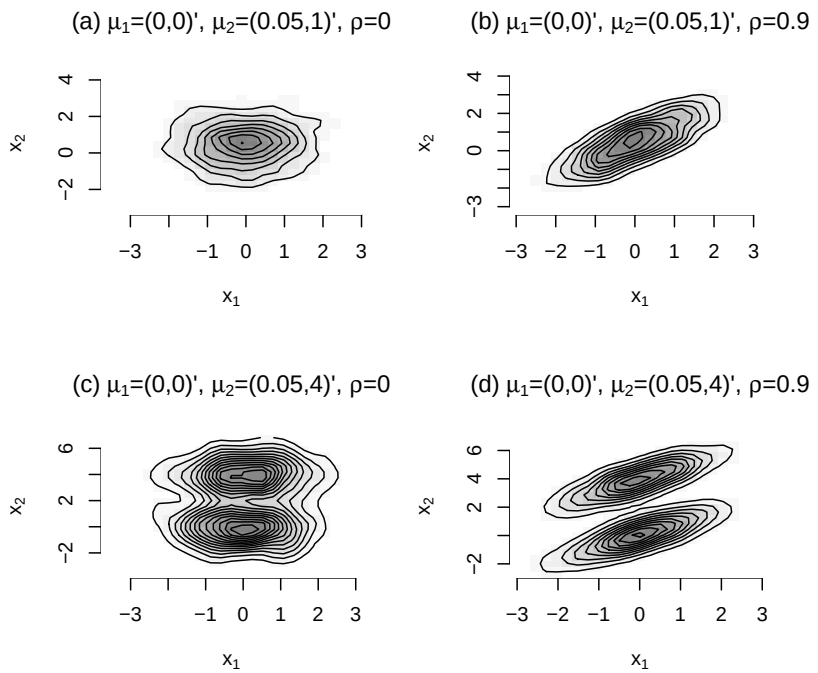

Figure 1: Empirical contour plots from 1000 samples simulated from density (3), all with $p = 0.5$, $\boldsymbol{V}_1 = \boldsymbol{V}_2$ and $\sigma_1 = \sigma_2 = 1$.

We first introduce some notations before presenting the main results that underlies the intuition. Given a sample of independent and identically distributed (i.i.d.) random variables $\mathbf{x}_1, ..., \mathbf{x}_n$ from the distribution (2), we write the log likelihood as

$$l(\boldsymbol{\theta}) = \sum_{i}^{n} \log f(\mathbf{x}_i).$$

The analytical forms of the score function and the Hessian matrix for arbitrary (finite) number of Gaussian mixtures with arbitrary dimension of observations are derived in Boldea and Magnus (2009). To simplify the analytical discussion, we focus on the simple case where the proportion p and variance matrices $\boldsymbol{V}_k$ are known, thus the only unknown parameters are $\boldsymbol{\theta} = \boldsymbol{\mu} = (\boldsymbol{\mu}_1, \boldsymbol{\mu}_2)'$.

Denote the score function by $\boldsymbol{q}(\boldsymbol{\mu}) = \sum_i \boldsymbol{q}_i(\boldsymbol{\mu})$, where

$$\boldsymbol{q}_i(\boldsymbol{\mu}) = \frac{\partial \log f(\mathbf{x}_i)}{\partial \boldsymbol{\mu}} = (\boldsymbol{q}_{i1}, \boldsymbol{q}_{i2}),$$

and the Hessian matrix by $\boldsymbol{Q}(\boldsymbol{\mu}) = \sum_i \boldsymbol{Q}_i(\boldsymbol{\mu})$, where

$$\boldsymbol{Q}_i(\boldsymbol{\mu}) = \frac{\partial^2 \log f(\mathbf{x}_i)}{\partial \boldsymbol{\mu} \partial \boldsymbol{\mu}'} = \begin{pmatrix} \boldsymbol{Q}_{i11} & \boldsymbol{Q}_{i12} \\ \boldsymbol{Q}_{i21} & \boldsymbol{Q}_{i22} \end{pmatrix}.$$

The maximum likelihood estimate (MLE) of the parameters, $\hat{\boldsymbol{\mu}}$, can be obtained via the EM algorithm by finding the solution to the system of equations of setting the score functions to 0. With the EM algorithm, the missing mixture membership is augmented iteratively for each unit and the likelihood is maximized conditional on the augmentations. In likelihood-based approaches, the variance is usually estimated from the information matrix. If the model is correctly specified, the information matrix is defined by

$$\mathcal{I} = -\mathbb{E}(\boldsymbol{Q}) = \mathbb{E}(\boldsymbol{q}\boldsymbol{q}'),$$

where the second equality holds because of the second-order regularity of f. The asymptotic variance of the MLE of $\boldsymbol{\mu}$ is $\mathcal{I}^{-1}$.

2.2 Analytical results

Our main analytical results are obtained through investigating the allocation probability $-$ the probability of unit i being in the group k:

$$\Pr(i \in k \mid \mathbf{x}_i) = \frac{c f_k(\mathbf{x}_i)}{f(\mathbf{x}_i)} = \left\{ 1 + \frac{\bar{c}\, f_{\bar{k}}(\mathbf{x}_i)}{c\, f_k(\mathbf{x}_i)} \right\}^{-1}, \tag{4}$$

where

$$c = \begin{cases} p, & \text{for } k = 1 \\ 1-p, & \text{for } k = 2 \end{cases}, \quad \bar{c} = \begin{cases} 1-p, & \text{for } k = 1 \\ p, & \text{for } k = 2 \end{cases}, \quad \bar{k} = \begin{cases} 2, & \text{for } k = 1 \\ 1, & \text{for } k = 2 \end{cases}.$$

The allocation probability tends to 1 or 0 the better the mixture disentanglement is, while it tends to p or $1-p$ the worse the mixture disentanglement is.

We first investigate the properties of the key term $f_{\bar{k}}(\mathbf{x}_i)/f_k(\mathbf{x}_i)$ in (4) for the special case where the two components have the same variance covariance matrix (homoscedasticity), $\boldsymbol{V}_1 = \boldsymbol{V}_2 = \boldsymbol{V}$. To simplify discussion, we make the following transformations:

$$\boldsymbol{\mu}_1 = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \quad \boldsymbol{\mu}_2 = \mathbf{d} = \begin{pmatrix} d_1 = \mu_{12} - \mu_{11} \\ d_2 = \mu_{22} - \mu_{21} \end{pmatrix}, \quad \boldsymbol{V} = \begin{pmatrix} \sigma_1^2 & \rho\, \sigma_1 \sigma_2 \\ \rho\, \sigma_1 \sigma_2 & \sigma_2^2 \end{pmatrix}. \tag{5}$$

Then the term $f_{\bar{k}}(\mathbf{x}_i)/f_k(\mathbf{x}_i)$ can be expressed as:

$$f_{\bar{k}}(\mathbf{x}_i)/f_k(\mathbf{x}_i) = \exp\left\{ \mathbf{d}' \mathbf{V}^{-1} (\mathbf{x}_i - \boldsymbol{d}/2) \right\} = \exp\left\{ h(\mathbf{x}_i) \right\}.$$

When $\mathbf{x}_i$ belongs to component $k = 1$, $h(\mathbf{x}_i) \sim \mathrm{N}(-\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d}/2,\ \boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d})$, while when $\mathbf{x}_i$ belongs to component $k = 2$, $h(\mathbf{x}_i) \sim \mathrm{N}(\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d}/2,\ \boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d})$. Further, we can show that the probability of $h(\mathbf{x}_i) < 0$ when $\mathbf{x}_i$ comes from group $k = 1$, as well as the probability of $h(\mathbf{x}_i) > 0$ when $\mathbf{x}_i$ comes from group $k = 2$, increases with $\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d}$ as

$$\begin{aligned} \Pr(h(\mathbf{x}_i) < 0 \mid k = 1) &= \Phi\left(-\mathbb{E}(h(\mathbf{x}_i))\, \mathbb{V}(h(\mathbf{x}_i))^{-1/2} \right) = \Phi((\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d})^{1/2}/2), \\ \Pr(h(\mathbf{x}_i) > 0 \mid k = 2) &= \Phi\left(\mathbb{E}(h(\mathbf{x}_i))\, \mathbb{V}(h(\mathbf{x}_i))^{-1/2} \right) = \Phi((\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d})^{1/2}/2). \end{aligned}$$

By the Bayes rule, we have

$$\begin{aligned}\Pr(k=1 \mid h(\mathbf{x}_i)<0) &= \frac{p\ \Pr(h(\mathbf{x}_i)<0 \mid k=1)}{p\ \Pr(h(\mathbf{x}_i)<0 \mid k=1)+(1-p)\ \Pr(h(\mathbf{x}_i)<0 \mid k=2)} \\ &= \left[1+\frac{1-p}{p}\left\{\frac{1}{\Phi((\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d})^{1/2}/2)}-1\right\}\right]^{-1},\end{aligned}$$

and similar formula for $\Pr(k=2 \mid h(\mathbf{x}_i)>0)$.

These results illustrate the critical role of $\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d}$, essentially a standardized distance between the two components, in disentangling the components: As $\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d}$ increases from the minimum 0 to infinity, the probability $\Pr(k=1 \mid h(\mathbf{x}_i)<0)$ increases monotonically from its minimum p towards the maximum 1, where the former is equivalent to a random assignment of the group membership for each unit and the latter is equivalent to the case when the group membership is known. Writing out $\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d}$ using the parameters in (5),

$$\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d}=(d_1^2\sigma_2^2+d_2^2\sigma_1^2-2\rho\sigma_1\sigma_1 d_1 d_2)/\{(1-\rho^2)\sigma_1^2\sigma_2^2\},$$

we have the following results.

RESULT 1. *For a bivariate mixture model of two Gaussian components with equal and known variance-covariance matrices, given the parameterization in equations* (2) *and* (5):

(1) *For fixed values of* $(d_1,\sigma_1,\sigma_2,\rho)$, $\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d}$ *reaches its minimum at* $d_2=\rho d_1\sigma_2/\sigma_1$, *and the minimum is* d_1^2/σ_1^2, *which is the same value of* $\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d}$ *in the univariate mixture model.*

(2) *For fixed values of* $(d_1,d_2,\sigma_1,\sigma_2)$, $\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d}$ *reaches its minimum at two mutually exclusive values of* ρ: $d_2\sigma_1/d_1\sigma_2$ *or* $d_1\sigma_2/d_2\sigma_1$, *and the minimum is either* d_1^2/σ_1^2*(the same value of* $\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d}$ *in the univariate mixture model) or* d_2^2/σ_2^2 *(a value strictly greater than* d_1^2/σ_1^2*), respectively.*

(3) *For fixed values of* $(d_1, \sigma_1, \sigma_2, \rho)$*, the probability of allocating unit* i *to group* k *when unit* i *indeed belongs to component* k*,* $cf_k(\mathbf{x}_i)/f(\mathbf{x}_i)$*, increases with* d_2 *and*

$$\lim_{|d_2|\to\infty} \{c\ f_k(\mathbf{x}_i)/f(\mathbf{x}_i) \mid (i \in k)\} = 1,$$

while for fixed values of $(d_1, d_2, \sigma_1, \sigma_2)$*, the probability of allocating unit* i *to group* k *when unit* i *indeed belongs to component* k*,* $cf_k(\mathbf{x}_i)/f(\mathbf{x}_i)$*, increases with* ρ *and*

$$\lim_{|\rho|\to 1} \{c\ f_k(\mathbf{x}_i)/f(\mathbf{x}_i) \mid (i \in k)\} = 1.$$

PROOF. See Appendix 1.

Result 1 states that, under correct model specification, the standardized distance between the groups and thus the probability of correctly assigning group membership for each unit from a bivariate model is always greater than or equal to that from the corresponding marginal univariate model, and it increases with the distance between the group means of the second variable, and/or the conditional correlation between the variables within components.

The allocation probability is closely related to the information matrix. Specifically, the score function

$$\boldsymbol{q}_{ik}(\boldsymbol{\mu}) = \partial \log f(\mathbf{x}_i)/\partial \boldsymbol{\mu}_k = c\ \{f_k(\mathbf{x}_i)/f(\mathbf{x}_i)\}\ \boldsymbol{V}_k^{-1}(\mathbf{x}_i - \boldsymbol{\mu}_k).$$

Consider the consistent estimator for the information matrix $-$ the outer product of the scores evaluated at the MLE:

$$\mathcal{I}_1 = \sum_i^n\ \boldsymbol{q}_i(\hat{\boldsymbol{\mu}})\ \boldsymbol{q}_i(\hat{\boldsymbol{\mu}})'. \tag{6}$$

The following result can be proved.

RESULT 2. *For a bivariate mixture model of two Gaussian components with equal and known variance-covariance matrices (with parameterization in equations* (2) *and* (5)*), given correct*

model specification, fixed $(d_1, \sigma_1, \sigma_2, \rho)$ *and fixed sample size* n*,*

$$\lim_{|d_2|\to\infty} \mathcal{I}_1 = \begin{pmatrix} \mathbf{V}^{-1} \sum_{i\in k=1} (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_1)(\mathbf{x}_i - \hat{\boldsymbol{\mu}}_1)' \mathbf{V}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{V}^{-1} \sum_{i\in k=2} (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_2)(\mathbf{x}_i - \hat{\boldsymbol{\mu}}_2)' \mathbf{V}^{-1} \end{pmatrix},$$

where $\hat{\boldsymbol{\mu}}_k$ *is the MLE of* $\boldsymbol{\mu}_k$*, and the diagonal blocks are the outer products of the scores for* $\boldsymbol{\mu}_k$ *when the mixture membership for each unit is known. The same result holds when* $|\rho| \to 1$*, for fixed* $(d_1, d_2, \sigma_1, \sigma_2)$ *and fixed sample size* n*.*

PROOF. See Appendix 2.

Distinct from standard asymptotic results regarding increasing sample size, Result 2 is obtained with fixed n but increasing values of d_2 or ρ. It implies that as the distance between the means of the secondary variable in two components or/and the conditional correlation between the two variables increases, the information number for the means of the primary variable converges to its maximum value − the one from an analysis with the component labels known.

Intuitively, similar results also hold for mixtures with unequal variance-covariance matrices. However, general analytical results are difficult to obtain. We consider a second special case, where the two variables are conditionally independent in each group, that is, $\rho_1 = \rho_2 = 0$, regardless of whether the variances σ_1, σ_2 are the same. Corresponding to Result 1 and 2, we have the following results.

RESULT 3. *For a bivariate mixture model of two Gaussian components with known variance-covariance matrices and* $\rho_1 = \rho_2 = 0$*, given correct model specification, fixed values of* $(d_1, \sigma_{11}, \sigma_{21}, \sigma_{12}, \sigma_{22})$ *and fixed sample size* n*:*

(1) *The probability of allocating unit* i *to group* k *when unit* i *indeed belongs to component* k*,* $cf_k(\mathbf{x}_i)/f(\mathbf{x}_i)$*, increases with* d_2 *and*

$$\lim_{|d_2|\to\infty} \{c\, f_k(\mathbf{x}_i)/f(\mathbf{x}_i) \mid (i \in k)\} = 1.$$

(2) *The estimated information matrix* $\mathcal{I}_1$

$$\lim_{|d_2|\to\infty} \mathcal{I}_1 = \begin{pmatrix} \mathbf{V}_1^{-1} \sum_{i\in k=1} (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_1)(\mathbf{x}_i - \hat{\boldsymbol{\mu}}_1)' \mathbf{V}_1^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{V}_2^{-1} \sum_{i\in k=2} (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_2)(\mathbf{x}_i - \hat{\boldsymbol{\mu}}_2)' \mathbf{V}_2^{-1} \end{pmatrix},$$

where $\hat{\boldsymbol{\mu}}_k$ *is the MLE of* $\boldsymbol{\mu}_k$*, and the diagonal blocks are the outer products of the scores for* $\boldsymbol{\mu}_k$ *when the mixture membership for each unit is known.*

PROOF. See Appendix 3.

These results are intuitive, because as the secondary outcome distribution is increasingly separated between the two components (increasing d_2), the component labels become clearer until completely known, regardless of whether the primary outcome distribution is well separated. A cross-dimensional comparison of the information number with fixed d_2, on the other hand, may be informative in practice, but is also more difficult to obtain. Below we present a result derived with fixed d_2 for the special case of equal $\boldsymbol{V}$ and $\rho = 0$.

RESULT 4. *For a bivariate mixture model of two Gaussian components with equal and known variance-covariance matrices and* $\rho = 0$*, given correct model specification and fixed values of* $(d_1, d_2, \sigma_1, \sigma_2)$*, the information numbers for the means of the primary variable in* $\mathcal{I}_1$ *are larger than the corresponding ones from the univariate model for a large sample size* n*.*

PROOF. See Appendix 4.

Result 4 is not a direct comparison of the estimated standard errors. Nevertheless, given $\rho = 0$, results from simulations show the off-diagonal terms of $\mathcal{I}_1$ from the bivariate model quickly disappear with increasing d_2. Consequently the estimated standard errors for the means of the primary variable from the bivariate model can be approximated by the inverse of their information numbers, which can be easily shown to be lower than the estimated standard error from the corresponding marginal univariate model, given the positive definiteness of covariances matrices.

The above results are established assuming correct model specification. However, the information gain from utilizing secondary variables is obtained at the cost of having to specify more complex multivariate models. The number of parameters to be estimated in mixture models increases rapidly with the number of variables involved in the anaysis, increasing model uncertainty and also the possibility of misspecification. In particular, multivariate normality is a much stronger assumption than univariate normality. It is therefore crucial to assess model assumptions in multivariate analysis. In the case of Gaussian mixtures, one way to assess normality and homoscedasticity is to apply the test of Hawkins (1981) to the clusters implied by the MLE. More discussions on this can be found in McLachlan and Basford (1988), Section 3.2, and McLachlan (1992), Chapter 6.

Another benefit of introducing an auxiliary variable, which will be partly shown in the empirical analysis below, is that it tends to regularize the model. Mixture models with Gaussian (as well as most uni-modal distributions) components are not regular in the sense the ML regularity conditions for the likelihood function only hold locally, so that the likelihood function will generally have multiple roots, only one of which corresponds to the efficient likelihood estimator. The prevalence and the likelihood of spurious roots tends to disappear with the introduction of an auxiliary variable that it is highly associated with the mixture membership. For the means of the primary outcome of a mixture of two Gaussians, the gain is intuitive: upon inspecting the general formula of the observed information numbers calculated as the outer product of gradients. These are linear combinations of the squared scores and, for each component, the introduction of an auxiliary variable tends to annul the addends provided by the units belonging to the wrong component, and taking only those from the correct component. Discarding the wrong information results in an observed information matrix having the structure of a diagonal block matrix with null off-diagonal blocks and diagonal blocks equal to those of two regular Gaussian models. For a very entangled mixture of two Gaussian distributions

the benefit in the standard error tends to be at most equivalent to multiply the sample size by the inverse of the mixture proportion, i.e., equivalent to doubling the sample size if the mixture proportion is equal to 0.5.

3 Simulations

General analytical results of cross-dimensional comparison for mixture models with arbitrary number of components and dimensions are difficult to establish; several analytical results were however obtained for some special cases. We conduct now simulation studies to investigate the small-sample behavior of bivariate analysis and the corresponding marginal analysis of mixture models under a wider range of settings. Specifically, we examine the estimated standard error of the MLE for the component means of the first variable μ_{1k} ($k = 1, 2$).

Besides $\mathcal{I}_1$, we also consider the estimator of the variance matrix based on the Hessian matrix of the likelihood:

$$\mathcal{I}_2 = -\boldsymbol{Q}(\hat{\boldsymbol{\mu}}) = -\sum_i \boldsymbol{Q}_i(\hat{\boldsymbol{\mu}}). \tag{7}$$

If the model is correctly specified, the inverses $\mathcal{I}_1^{-1}$ and $\mathcal{I}_2^{-1}$ are both consistent estimators of the asymptotic variance of $\hat{\boldsymbol{\mu}}$. The closed-form Hessian matrix for $\boldsymbol{\mu}$ was derived in Boldea and Magnus (2009). Under model mis-specification, we will also consider the robust "sandwich" estimator (Huber, 1967):

$$\mathcal{I}_3^{-1} = \hat{\mathbb{V}}(\hat{\boldsymbol{\mu}}) = \mathcal{I}_2^{-1}\mathcal{I}_1\mathcal{I}_2^{-1}, \tag{8}$$

which is a consistent estimator for the variance, whether or not the model is correctly specified.

We consider three simulation settings, all with the sample size $n = 500$ and the weight of component 1, $p = 0.4$.

S1. *Correctly specified model with known covariance matrices.* The data is generated from the bivariate Gaussian mixture density (2)-(3), with $\boldsymbol{\mu}_1 = (0, 0)'$, $\boldsymbol{\mu}_2 = (1, d_2)'$, and

$\boldsymbol{V}_1 = \boldsymbol{V}_2 = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$. The parameters are estimated assuming the Gaussian mixture structure with p and $\boldsymbol{V}$ fixed at the true values.

S2. *Correctly specified model with unknown covariance matrices*. The data is generated from the bivariate Gaussian mixture density (2)-(3) with the same true parameters as in S1. The parameters are estimated assuming the Gaussian mixture structure with unknown p and $\boldsymbol{V}$.

S3. *Misspecified models (skewed with heavy tail) with unknown covariance matrices*. The data is generated from a mixture of two bivariate non-central *t* distributions whose marginals have the same shape, using the following steps:

(i) Draw a sample of size $n = 500$ from a Bernoulli distribution with $\Pr(i = 1) = 0.4$ and let n_1 be the number of time that $i = 1$ and $n_2 = n - n_1$.

(ii) For $i = 1, 2$ draw $2n_i$ random number from the univariate non-central t distribution with degree of freedom $df = 20$ and non-centrality parameter $\lambda = 7$ using the formula: $t = (Z + \lambda)/(W/df)^{1/2}$, where $Z \sim \mathrm{N}(0, 1)$, $W \sim \chi^2_{df}$.

(iii) Standardize these random draws to have mean 0 and variance 1 (the mean of the non-central *t* with $df = 20$ and $\lambda = 7$ is: $\lambda \left(\frac{df}{2}\right)^{1/2} \Gamma(\frac{df-1}{2})/\Gamma(\frac{df}{2}) \approx 7.28$, and the variance is $\frac{df\ (1+\lambda^2)}{df-2} - \frac{\lambda^2\ df}{2} \left\{\Gamma(\frac{df-1}{2})/\Gamma(\frac{df}{2})\right\}^2 \approx 2.60$), and arrange the standardized numbers in bivariate vectors $\epsilon_{i,1}, ..., \epsilon_{i,n_i}$

(iv) Transform $\epsilon_{i,j}$ to $\mathbf{x}_{i,j}$ (for $j = 1, ..., n_i$) by $\mathbf{x}_{i,j} = \boldsymbol{\mu}_i + \epsilon_{i,j}\boldsymbol{C}_i$, with $\boldsymbol{\mu}_i = (\mathbb{E}(x_{1i}), \mathbb{E}(x_{2i}))'$ and $\boldsymbol{C}_i$ being the Choleski decomposition of the desired correlation matrix $\boldsymbol{V} = \boldsymbol{C}_i\boldsymbol{C}_i'$.

It is straightforward to show the above steps simulate the set $\{\mathbf{x}_{i,j}\}$ that satisfies $\mathbb{E}(x_{11}) = 0, \mathbb{E}(x_{12}) = 1, \mathbb{E}(x_{21}) = 0$, $\rho(x_{1k}, x_{2k}) = \rho$ for $k = 1, 2$, and $\mathbb{V}(x_{mk}) = 1$ for $m, k = 1, 2$.

For each setting, we conduct two series of simulations: (1) fixing $\rho = 0$ and increasing the difference between the means of the second variable d_2; (2) fixing $d_2 = 0$ and increase the correlation between the two variables within each component ρ. By Result 1.1, $d_2 = 0$ leads to the smallest allocation probability when $\rho = 0$, while by Result 1.2, $\rho = 0$ leads to the smallest allocation probability when $d_2 = 0$.

The MLEs of the parameters are obtained from the EM algorithm, and the standard errors of the MLE are estimated from $\mathcal{I}_1$, $\mathcal{I}_2$ and $\mathcal{I}_3$. Following Boldea and Magnus (2009), for each setting we obtain the Monte Carlo (MC) approximation to the true standard error of $\hat{\mu}_{11}$ (and $\hat{\mu}_{12}$) as the standard deviations of the empirical distributions of the MLE $\hat{\mu}_{11}$ and $\hat{\mu}_{12}$ from $R = 10000$ replicates, each of sample size $n = 500$. The subsequent estimated standard errors, from $\mathcal{I}_1$, $\mathcal{I}_2$, other than $\mathcal{I}_3$ for S3, are assessed in terms of bias and root mean square error (RMSE) to the "true standard errors", calculated from 1000 replicates.

The estimated standard errors of the MLE for $\hat{\mu}_{11}$ from estimating a bivariate normal mixture models versus those from estimating the marginal univariate model are summarized in Table 1, 2, 3 under settings S1, S2, S3, respectively. The last column of each table reports the estimated allocation rate, which is an estimate of the proportion of units that are correctly allocated to the components. The allocation rate is a useful indicator for quantifying mixture disentanglement; it is here estimated by averaging the higher probability of unit i being in the group k calculated at the MLE (McLachlan and Basford, 1988): $\{\sum_i \max_k \Pr(i \in k \mid \mathbf{x}_i)\}/n$. The lower bound for the estimated allocation rate is 0.5; low values correspond to poor mixture disentanglements, and vice versa.

When the model is correctly specified with known variance (setting S1), as predicted by the analytical results, the bivariate analysis nearly always outperforms the univariate analysis, and the improvement increases as the distance between the two mixture components of the secondary variable or the correlation between the two variables increases. Here the estimator

		$\mathcal{I}_1$				$\mathcal{I}_2$				
	d_2 or ρ	|bias|	mean $\widehat{\text{se}}$	RMSE	(*)	|bias|	mean $\widehat{\text{se}}$	RMSE	(*)	$\widehat{\text{AR}}$
univ	-	$1.8e^{-3}$	$1.1e^{-1}$	$1.5e^{-2}$	-	$2.5e^{-3}$	$1.1e^{-1}$	$1.3e^{-2}$	-	.703
	0	$3.4e^{-3}$	$1.1e^{-1}$	$1.3e^{-2}$	345	$3.2e^{-3}$	$1.1e^{-1}$	$1.5e^{-2}$	531	.694
biv	1	$1.1e^{-3}$	$1.0e^{-1}$	$6.2e^{-3}$	993	$1.6e^{-3}$	$1.0e^{-1}$	$5.7e^{-3}$	911	.769
($\rho = 0$)	3	$1.9e^{-4}$	$7.7e^{-2}$	$4.3e^{-3}$	998	$8.9e^{-4}$	$7.6e^{-2}$	$2.4e^{-3}$	1000	.945
	5	$1.2e^{-4}$	$7.2e^{-2}$	$4.0e^{-3}$	1000	$7.2e^{-4}$	$7.2e^{-2}$	$2.0e^{-3}$	1000	.996
	50	$6.5e^{-4}$	$7.1e^{-2}$	$4.0e^{-3}$	999	$6.6e^{-5}$	$7.1e^{-2}$	$1.8e^{-3}$	1000	1.00
	.50	$9.6e^{-4}$	$1.1e^{-1}$	$9.7e^{-3}$	839	$1.8e^{-3}$	$1.1e^{-1}$	$9.8e^{-3}$	740	.726
biv	.75	$1.6e^{-3}$	$9.8e^{-2}$	$5.9e^{-3}$	999	$8.6e^{-4}$	$9.8e^{-2}$	$4.7e^{-3}$	948	.783
($d_2 = 0$)	.90	$4.4e^{-4}$	$8.5e^{-2}$	$4.7e^{-3}$	1000	$2.5e^{-4}$	$8.4e^{-2}$	$2.6e^{-3}$	1000	.878
	.99	$6.3e^{-4}$	$7.1e^{-2}$	$4.3e^{-3}$	999	$6.0e^{-5}$	$7.1e^{-2}$	$1.9e^{-3}$	1000	1.00

Table 1: The estimated standard error of the MLE of μ_{11} from bivariate analysis and corresponding marginal analysis under the simulation setting S1. The upper panel is with fixed $\rho = 0$ and varying d_2, and the lower panel is with fixed $d_2 = 0$ and varying ρ. Here "|bias|" stands for absolute bias; "mean $\widehat{\text{se}}$" stands for mean of the estimated s.e. of $\hat{\mu}_{11}$, (*) is number of times that, over the 1000 replications, the bivariate estimated s.e. of $\hat{\mu}_{11}$ is smaller than the univariate one, and "$\widehat{\text{AR}}$" is the estimated allocation rate.

$\mathcal{I}_2$ leads to comparable standard errors but smaller bias (thus smaller MSE) than $\mathcal{I}_1$. Interestingly, the ratio between the bivariate and the univariate mean $\widehat{\text{se}}$ has a lower bound of size $7.1e^{-2}/1.1e^{-1} \cong \sqrt{0.4}$ i.e. $\frac{\sqrt{n}}{\sqrt{n/p}} = \sqrt{p}$. Thus, the reduction of the s.e. is equivalent to a reduction due to increase the sample size by the inverse of the mixture proportion. Results for the estimated s.e. of the MLE of $\hat{\mu}_{12}$ and for alternative sample sizes, not reported here, confirm this evidence.

		$\mathcal{I}_1$				$\mathcal{I}_2$				
	d_2 or ρ	|bias|	mean $\widehat{\text{se}}$	RMSE	(*)	|bias|	mean $\widehat{\text{se}}$	RMSE	(*)	$\widehat{\text{AR}}$
univ	-	$1.7e^{-3}$	$1.1e^{-1}$	$1.5e^{-2}$	-	$2.5e^{-3}$	$1.1e^{-1}$	$1.3e^{-2}$	-	.893
	0	$3.4e^{-1}$	$2.8e^{-1}$	$4.1e^{-1}$	621	$3.4e^{-1}$	$2.8e^{-1}$	$6.3e^{-1}$	573	.862
biv	1	$2.8e^{-1}$	$3.0e^{-1}$	$4.0e^{-1}$	587	$3.2e^{-1}$	$2.5e^{-1}$	$4.0e^{-1}$	557	.871
($\rho = 0$)	3	$2.1e^{-3}$	$8.8e^{-2}$	$9.7e^{-3}$	804	$2.8e^{-3}$	$9.3e^{-2}$	$4.8e^{-2}$	782	.946
	5	$2.4e^{-4}$	$7.2e^{-2}$	$4.6e^{-3}$	839	$3.0e^{-3}$	$7.5e^{-2}$	$4.0e^{-2}$	819	.995
	50	$6.4e^{-4}$	$7.2e^{-2}$	$4.3e^{-3}$	823	$1.1e^{-3}$	$7.2e^{-2}$	$1.6e^{-2}$	807	1.00
	.50	$3.1e^{-1}$	$3.0e^{-1}$	$4.1e^{-1}$	606	$3.4e^{-1}$	$2.7e^{-1}$	$4.8e^{-1}$	558	.863
biv	.75	$2.4e^{-1}$	$2.9e^{-1}$	$4.5e^{-1}$	573	$2.7e^{-1}$	$2.6e^{-1}$	$3.5e^{-1}$	533	.867
($d_2 = 0$)	.90	$4.9e^{-2}$	$1.6e^{-1}$	$7.1e^{-2}$	670	$4.2e^{-2}$	$1.7e^{-1}$	$1.1e^{-1}$	630	.887
	.99	$6.5e^{-4}$	$7.2e^{-2}$	$4.4e^{-3}$	846	$2.6e^{-3}$	$7.3e^{-2}$	$6.2e^{-2}$	839	1.00

Table 2: The estimated standard error of the MLE of μ_{11} from bivariate analysis and corresponding marginal analysis under the simulation setting S2. The upper panel is with fixed $\rho = 0$ and varying d_2, and the lower panel is with fixed $d_2 = 0$ and varying ρ.

When the model is correctly specified but with unknown variance (setting S2), the bivariate analysis still leads to smaller standard errors than the univariate analysis in at least 60% of the time, and this rate increases to over 80% as d_2 increases to 3. But unlike in setting S1, here $\mathcal{I}_2$ leads to comparable standard error but larger bias (thus larger MSE) than $\mathcal{I}_1$. Interestingly, in both settings S1 and S2, the improvement in bias and MSE of bivariate analysis appears to plateau after d_2 reaches 5 despite the estimated allocation rate continuing to increase with d_2. This illustrates that, in practice, a secondary variable with even modest distance between the two components is sufficent to provide noticeable improvement.

For setting S3, the "sandwich" estimator $\mathcal{I}_3$ yields standard errors for the MLE that are robust to specification error. However, it ignores bias, which may be appreciable, so that results

		$\mathcal{I}_1$			$\mathcal{I}_2$			$\mathcal{I}_3$			
	d_2 or ρ	|bias|	mean $\widehat{\text{se}}$	RMSE	|bias|	mean $\widehat{\text{se}}$	RMSE	|bias|	mean $\widehat{\text{se}}$	RMSE	$\widehat{\text{AR}}$
	4	$4.2e^{-3}$	$8.1e^{-2}$	$9.5e^{-3}$	$9.6e^{-4}$	$7.6e^{-2}$	$1.9e^{-2}$	$1.4e^{-2}$	$9.0e^{-2}$	$6.6e^{-2}$	.986
$\rho = 0$	5	$7.0e^{-3}$	$7.9e^{-2}$	$1.1e^{-2}$	$2.1e^{-3}$	$7.4e^{-2}$	$1.9e^{-2}$	$1.7e^{-2}$	$8.9e^{-2}$	$7.3e^{-2}$	.996
	50	$7.3e^{-3}$	$7.8e^{-2}$	$1.1e^{-2}$	$2.1e^{-4}$	$7.1e^{-2}$	$9.8e^{-3}$	$2.5e^{-3}$	$7.4e^{-2}$	$2.3e^{-2}$	1.00
	0.50	$2.5e^{-3}$	$8.0e^{-2}$	$9.0e^{-3}$	$4.4e^{-3}$	$7.3e^{-2}$	$1.8e^{-2}$	$1.7e^{-3}$	$8.0e^{-2}$	$4.7e^{-2}$	.987
$d_2 = 4$	0.75	$6.5e^{-3}$	$7.9e^{-2}$	$1.1e^{-2}$	$9.6e^{-4}$	$7.2e^{-2}$	$2.0e^{-2}$	$6.0e^{-3}$	$7.9e^{-2}$	$6.0e^{-2}$	.995
	0.90	$7.1e^{-3}$	$7.8e^{-2}$	$1.0e^{-2}$	$1.5e^{-3}$	$7.3e^{-2}$	$2.1e^{-2}$	$1.0e^{-2}$	$8.1e^{-2}$	$6.9e^{-2}$	.999
	0.99	$6.6e^{-3}$	$7.8e^{-2}$	$9.9e^{-3}$	$1.3e^{-3}$	$7.1e^{-2}$	$1.6e^{-2}$	$1.7e^{-3}$	$7.3e^{-2}$	$4.7e^{-2}$	1.00

Table 3: The estimated standard error of the MLE of μ_{11} from bivariate analysis, under the simulation setting S3. The upper panel is with fixed $\rho = 0$ and varying d_2, and the lower panel is with fixed $d_2 = 4$ and varying ρ.

can be misleading (e.g. Freedman, 2006). Consequently, we have considered only values of d_2 and ρ such that the pseudo-MLE provides a good approximation to the true data model: the assessments have been carried out when the bias $\bar{\hat{\mu}}_{11} - \mathbb{E}(X_{11})$ is low, namely less than 0.03 (with $\bar{\hat{\mu}}_{11}$ be the average of the empirical distribution of the MLE of $\hat{\mu}_{11}$ over the 10000 replicates). As a result, we do not assess the performance of bivariate versus univariate estimators this time, given the bad pseudo-MLE approximation obtained for the latter (the bias over the 10000 replicates is $\bar{\hat{\mu}}_{11} - \mathbb{E}(X_{11}) = 0.305$). Moreover, the analysis for increasing values of ρ has been carried out by fixing the distance d_2 to 4 instead of 0, since the posing of d_2 to 0 would have resulted in a good approximation of the pseudo-MLE only when $\rho = 0.999$. Table 3 shows the outer product estimator $\mathcal{I}_1$ leads to smaller standard error and comparable bias, thus a better performance in term of RMSE, than both $\mathcal{I}_2$ and the sandwich estimator $\mathcal{I}_3$. The extra advantage of the bivariate analysis when the underlying model is incorrect is in the great

reduction of bias, compared to the univariate case, leading to a really robust inference (more pronounced for increasing d_2 than for increasing ρ). The large value of the bias obtained for the univariate analysis shows it fails to provide a good approximation to the true data model, and signals an analysis of resulting MLE of μ_{11} would be misleading.

4 Real applications

4.1 Crab data

The crab data of the genus *Leptograpsus variegatus*, originally collected by Campbell and Mahon (1974), has been often analyzed in the literature of multivariate mixture models (e.g. Ripley, 1996; McLachlan and Peel, 1998, 2000). Here we focus the sample of $n = 100$ blue crabs, with $n_1 = 50$ males and $n_2 = 50$ females, corresponding to the two components with component labels known. Each specimen has measurements (in mm) on the width of the front lip (FL), the rear width (RW), the length along the midline (CL), the body depth (BD), and the carapace width (CW). We use the data to conduct cross-dimensional comparison of the mixture models with more than two variables. While Hawkins' test suggests both normality and homoscedasticity assumptions to be reasonable here, McLachlan and Peel (2000) found that homoscedasticity may lead to inferior model fitting. For illustration purpose, we consider the hypothetical setting that RW is of primary interest and all other variables are secondary. We performed three clustering analyses, ignoring the known component labels: In the first, we fitted a univariate mixture model to RW, in the second we fitted a bivariate model to RW and CL, and in the third we fitted trivariate models to RW and CL with either FL, or BD, or CW as an additional third variable; all the models were with two Gaussian components and heterogeneous covariance matrices. The MLEs of the parameters were obtained running the EM algorithm with several random starting values. The labelling of mixtures components are

those obtained by setting the starting values in both analyses as the component-specific sample means (e.g., 11.72 and 12.14 for RW; 32.01 and 28.10 for CL), variances (e.g., 4.46 and 5.95 for RW; 53.42 and 35.04 for CL) and covariances. The results are reported in Table 4. In the univariate analysis, the EM converged to a spurious maximum point, with $\hat{\sigma}^2_{\text{RWmales}} = .68$, resulting from a group of eight outliers being erroneously identified as a component. As a consequence, the A.R. was very low (.10 for the males and .50 overall). The bivariate analysis reduced the adverse effect of outliers, leading to $\hat{\sigma}_{\text{RWmales}} = 2.82$ and the overall A.R. improves from .50 to .87. Also the standard errors from all three estimators improved for the males (the comparison for females is not meaningful due to the spurious point). The significant improvement is as expected from the theoretical results because the empirical correlation (given the true labels) between RW and CL is very high for both males (.977) and females (.987). The trivariate analyses lead to comparable results as the bivariate one. This plateau in performance is not surprising since the information gain from adding variables is obtained at the price of the extra uncertainty in estimating more parameters; the latter can outweigh the former especially when a lower-dimensional analysis already produces accurate results.

Besides RW, we have also run similar analysis with CL as the primary outcome. No spurious point was detected. The standard errors of the cluster means estimated from all the three estimators reduced significantly (60% in males and 80% in females) from the univariate to the bivariate analysis and the overall A.R. increased from .60 to .87. Same as before, the trivariate analyses did not provide further improvement (details are omitted here).

In the presence of multiple candidate secondary variables, we suggest the following selection procedure: first, conduct a normality test (e.g., Hawkins') for each bivariate pair of the primary variable and one secondary variable and select the ones deemed normal; second, for each selected pair perform a bivariate mixture analysis and, given the estimated labels, calculate the empirical within-component correlations and the distance between the component

		MLE		s.e. for $\hat{\mu}_{\text{RW}}$			A.R.	
		$\hat{\mu}_{\text{RW}}$	$\hat{\sigma}^2_{\text{RW}}$	$\mathcal{I}_1$	$\mathcal{I}_2$	$\mathcal{I}_3$	component	overall
male	univ	7.97	0.68	.83	.65	.53	.100	.500
	biv (CL)	12.40	2.82	.36	.38	.35	.740	.870
	triv (CL,FL)	12.43	2.72	.34	.32	.32	.740	.870
	triv (CL,CW)	12.44	2.73	.43	.32	.39	.720	.860
	triv (CL,BD)	12.32	2.81	.38	.33	.35	.800	.900
female	univ	12.27	4.04	.36	.31	.29	.900	.500
	biv (CL)	11.63	6.40	.34	.34	.34	1.00	.870
	triv (CL,FL)	11.61	6.41	.39	.33	.34	1.00	.870
	triv (CL,CW)	11.61	6.39	.36	.33	.37	1.00	.860
	triv (CL,BD)	11.66	6.58	.39	.35	.34	1.00	.900

Table 4: The estimated mean and variance of RW, and the standard error of the mean for male and female crabs. In the bivariate model, CL is as the second variable; in the trivariate analyses, FL, CW, BD are separately used as the third variable besides RW and CL. The sample mean of RW for the males and females is 11.72 and 12.14, respectively, and the sample variance of RW for the males and females is 4.46 and 5.95, respectively.

means of the secondary variable; third, choose the secondary variable that gives the highest absolution correlations or/and distances.

4.2 Educational cost of World War II

The second application arises from the Instrumental Variable (IV) approach in causal inference (Angrist et al., 1996), which inherently defines a mixture structure as shown later. An instru-

mental variable or instrument is a variable that is correlated with the treatment variable, but does not have a direct effect on the outcome, only indirectly through the treatment variable. The instrumental variable is often viewed as defining a natural experiment. Ichino and Winter-Ebmer (2004) used the IV approach to evaluate the long-run educational effect of World War II on earnings. In particular, they used the cohort of birth as an instrument (Z): $Z = 1$ for individuals born between 1930 and 1939 (these individuals were in primary school age during the war and the immediately following period) and $Z = 0$ otherwise. It is reasonable to assume that which year an individual was born is random (by nature) and does not directly affect one's earnings later in life once accounting for the secular trend towards higher earnings, but it can indeed affect the education level an individual received ($D = 1$ for poorly educated, $D = 0$ otherwise) due to the intervention of war, which in turn affects the earnings later. The population can be divided into four latent subpopulations according to an individual's potential (counterfactual) educational levels under different values of the instrument:

1. *Always-poorly educated* ($G = a$): individuals who would obtain low education levels irrespective of the cohort of birth;
2. *Never-poorly educated* ($G = n$): individuals who would obtain high education levels irrespective of the cohort of birth;
3. *Compliers* ($G = c$): individuals who would obtain low education level if born in the decade immediately before war, but would obtain high education level if not born in the decade immediately before war;
4. *Defiers* ($G = d$), individuals who would obtain high education level if born in the decade immediately before war but low education level otherwise.

It is usually reasonable to rule out defiers in practice. And the estimand of interest lies in the effect of war on earnings for compliers, known as the compliers average causal effect (CACE).

The inferential challenge is that the individual's subclass membership is not always observed. Specifically, given the structural assumptions in Angrist et al. (1996), individuals for which we observe $(Z = 0, D = 1)$ are never-poorly educated, and those for which $(Z = 0, D = 1)$ are always-poorly educated. But individuals with $(Z = 1, D = 1)$ or $(Z = 0, D = 0)$ consists of two different mixtures, that are a mixture of always-poorly educated and compliers, and a mixture of never-poorly educated and compliers respectively, as shown in Table 5. For continuous outcomes, Imbens and Rubin (1997) proposed the following mixture model for the above formulation of instrumental variable:

$$\begin{aligned} f(x,d,z) &= \mathbf{1}_{\{D=1,Z=0\}} \cdot (1-\pi) \cdot \omega_a \cdot \phi_{a0}(x) + \mathbf{1}_{\{D=0,Z=1\}} \cdot \pi \cdot \omega_n \cdot \phi_{n1}(x) \\ &\quad + \mathbf{1}_{\{D=1,Z=1\}} \cdot \pi \cdot \{\omega_a \cdot \phi_{a1}(x) + \omega_c \cdot \phi_{c1}(x)\} \\ &\quad + \mathbf{1}_{\{D=0,Z=0\}} \cdot (1-\pi) \cdot \{\omega_n \cdot \phi_{n0}(x) + \omega_c \cdot \phi_{c0}(x)\}, \end{aligned} \tag{9}$$

where $f(x,d,z)$ is the density of the outcome x given the observed instrument z and education level d, $\mathbf{1}_{\{A\}}$ is an indicator function of set A, π is the probability $P(Z=1)$, ω_g is the mixing probability, which is the probability of an individual being in the g group $P(G=g)$ for $g = a, c, n$, and $\phi_{gz}(x) = \phi_{gz}(x; \mu_{gz}, \sigma_{gz})$ is the (normal) outcome distribution for a unit in the g group that is assigned to the treatment z. The mixture structure is clearly shown in two last factors in (9), which are linked to each other by means of the two parameters π and ω_c, and therefore two separate analyses of the mixtures would lead to different results compared to the joint analysis of (9).

Our analysis uses the same data of Ichino and Winter-Ebmer (2004), collected from the wave 1986 of the German Socio-Economic Panel from which we consider only males born between 1925 and 1949. They defined the long-run educational cost of World War II on earnings in Germany as the average earnings loss experienced by those individuals who received less education because they were about in primary school age during the war or the immediately

	$Z = 0$	$Z = 1$
$D = 0$	never-poorly educated and compliers	never-poorly educated
$D = 1$	always-poorly educated	always-poorly educated and compliers

Table 5: Composition of the population classified by D and Z.

following years. To account for increasing trend of earnings with respect to age, following these authors, the primary outcome X is defined as the residual of a regression of natural log of average hourly earnings observed in 1986 in Germany on a cubic polynomial in age. The log transformation for income is usually adopted in the labor economics in order to induce normality in such otherwise asymmetric variable. To account for decreasing trend of earnings with respect to age, the "treatment" D is defined to be equal to one if the individual's residual of a regression of years of education on a cubic polynomial in age is smaller than the residuals' sample average (poorly educated) and zero otherwise. The instrument Z was defined earlier. In addition, we select an auxiliary variable: the hours worked per week. We will compare the results from fitting the univariate version of model (9) to earnings, versus those obtained from fitting the bivariate version of model (9) to earnings and the auxiliary variable.

We first eliminate the multivariate outliers detected on the initial sample of 1163 units. A subsequent visual checking of the histograms of the auxiliary variable for the two non-mixtures factors in (9) reveals a deviation from normality due to a slight bimodality. Individuals presenting very high level of hours worked have been consequently eliminated so that the final sample size results in 993 units. As shown in Mercatanti (2013), despite model (9) is identified, the main problem associated with a likelihood analysis arises from the possibility of having multiple roots for the likelihood equations, which results from the two mixtures of distributions being involved. We adopt here the proposed solution to identify the Efficient Likelihood Estimator (ELE) for parameters in (9), under heteroscedastic conditions for the

mixtures, as the local maximum likelihood point closest to the method of moments estimate of the mixing probabilities.

Table 6 reports the results of the univariate versus the bivariate analysis. The parameters related to the groups of units for which the subclass membership is observed (units with $D = 1$ and $Z = 0$, or $D = 0$ and $Z = 1$) do not involve mixture structure, and unsurprisingly their estimated values ($\hat{\mu}_{a0}$, $\hat{\sigma}_{a0}$, $\hat{\mu}_{n1}$, $\hat{\sigma}_{n1}$) and standard errors do not change from the univariate to the bivariate analysis. Significant reductions in the estimated standard errors have been obtained for the parameters of the mixture of always-poorly educated and compliers: this is a very entangled mixture for which the contribution of the secondary outcome is decisive to sharpen the inference. In particular for the group of compliers, for which the estimated standard errors show a reduction of about 55% in $\hat{\mu}_{c1}$, and 60% in $\hat{\sigma}_{c1}$. Moreover the most important estimand in this study, that is the average causal effect on earnings for compliers, $\hat{\mu}_{c1} - \hat{\mu}_{c0}$ (individuals whose educational choices were affected by the war), shows a reduction in its standard error ranging between 24%, by $\mathcal{I}_3$, to 32%, by $\mathcal{I}_2$, due essentially to the strong reduction observed for $\hat{\mu}_{c1}$. This leads to a decrease in the *p*-value for this quantity from about 0.700 for the univariate case to 0.054 ($\mathcal{I}_1$), 0.048 ($\mathcal{I}_2$), 0.081 ($\mathcal{I}_3$) for the bivariate case. The estimated standard errors for the rest of the parametric set show lighter reductions apart from the slight increases in $\hat{\sigma}_{c0}$ and $\hat{\mu}_{a1}$.

Interestingly, another advantage of the bivariate analysis in this example emerges from the analysis of the local maximum likelihood points detected. The ELEs reported in Table 6 correspond to the roots closest to the method of moments estimates of the mixing probabilities. The second closest root detected for the univariate case reports parameters values similar of that obtained for the ELE in the bivariate case (even if with generally larger estimated standard errors). The mixture composed by always-poorly educated and compliers is very entangled in the univariate case; this complicates the analysis so that this solution remains confused

with others local roots. The more effective disentanglement of the mixture allowed by the introduction of the secondary outcome succeeds in highlighting this solution as the ELE.

The practical interpretation of the results has to account for the definition of the outcome as log of earnings. This means the estimated effects are not differences in average amounts of money, but they are semi-elasticities, i.e. they show the approximate average percentage changes in earnings between groups of individuals classified by the cohort of birth. As expected the estimated effect for compliers is negative: the earnings are on average 25.73% lower for compliers who were affected by war because in primary school age during that period. The effect for always-poorly educated is substantially zero, while that for never-poorly educated results is positive – this is not surprising because it is reasonable to think that never-poorly educated individuals born between 1930 and 1939 took advantage of the lower average education level in their cohort by experiencing less competitive labour market conditions during their adulthood, thus increasing their average earnings.

5 Conclusion

We propose to sharpen the inference for a lower-dimensional mixture model by jointly modeling the primary variable and an auxiliary variable. We have established analytical results for several special cases that show that the probability of correctly allocating mixture memberships and the information number for the means of the primary outcome in a bivariate mixture model with two Gaussian components are generally larger than those in the corresponding univariate model. The improvement under more general settings, including misspecified models, is also observed in a comprehensive simulation study and in two real data analyses. As shown in the second empirical example, there is in general no need to include many auxiliary variables, as most of the information gain comes from the auxiliary variable with a high association with the

	Univariate case				Bivariate case			
	ELE	s.e.			ELE	s.e.		
		$\mathcal{I}_1$	$\mathcal{I}_2$	$\mathcal{I}_3$		$\mathcal{I}_1$	$\mathcal{I}_2$	$\mathcal{I}_3$
$\hat{\omega}_a$	0.7316	0.0293	0.0293	0.0293	0.7301	0.0293	0.0288	0.0290
$\hat{\omega}_n$	0.2075	0.0192	0.0193	0.0196	0.2099	0.0186	0.0180	0.0186
$\hat{\omega}_c$	0.0608	0.0239	0.0239	0.0246	0.0599	0.0230	0.0206	0.0220
$\hat{\mu}_{a0}$	-0.1229	0.0142	0.0142	0.0143	-0.1229	0.0142	0.0139	0.0143
$\hat{\mu}_{a1}$	-0.1333	0.0171	0.0174	0.0180	-0.1202	0.0185	0.0174	0.0194
$\hat{\mu}_{a1} - \hat{\mu}_{a0}$	-0.0104	0.0222	0.0228	0.0234	0.0027	0.0233	0.0225	0.0257
$\hat{\mu}_{n0}$	0.2724	0.0320	0.0309	0.0301	0.2585	0.0289	0.0280	0.0283
$\hat{\mu}_{n1}$	0.3104	0.0318	0.0318	0.0318	0.3104	0.0318	0.0318	0.0318
$\hat{\mu}_{n1} - \hat{\mu}_{n0}$	0.0380	0.0451	0.0443	0.0438	0.0519	0.0430	0.0423	0.0425
$\hat{\mu}_{c0}$	0.0351	0.1233	0.1233	0.1243	0.0685	0.1227	0.1160	0.1241
$\hat{\mu}_{c1}$	-0.0385	0.1398	0.1445	0.1507	-0.1888	0.0564	0.0564	0.0747
$\hat{\mu}_{c1} - \hat{\mu}_{c0}$	-0.0736	0.1909	0.1921	0.1951	-0.2573	0.1339	0.1305	0.1476
$\hat{\sigma}_{a0}$	0.2874	0.0082	0.0100	0.0121	0.2874	0.0082	0.0100	0.0121
$\hat{\sigma}_{a1}$	0.2472	0.0124	0.0128	0.0134	0.2796	0.0112	0.0117	0.0132
$\hat{\sigma}_{n0}$	0.2288	0.0237	0.0268	0.0305	0.2409	0.0214	0.0197	0.0196
$\hat{\sigma}_{n1}$	0.3034	0.0198	0.0225	0.0255	0.3034	0.0198	0.0225	0.0255
$\hat{\sigma}_{c0}$	0.4621	0.0856	0.0745	0.0652	0.4726	0.0900	0.0783	0.0695
$\hat{\sigma}_{c1}$	0.4551	0.1056	0.0892	0.0774	0.1138	0.0450	0.0309	0.0236

Table 6: Efficient Likelihood Estimates and Standard Errors from the univariate and the bivariate models applied to World War II data.

mixture membership.

The formal results we have obtained can be useful in many settings, such as those mentioned in the introduction. The goal in empirical analysis, e.g. in causal inference and small area estimation, should be to pick the *best* auxiliary variable that increases precision without increasing the risk of mis-specification. This issue will be the subject of our future investigations.

Acknowledgements

Mercatanti's research was partially supported by the U.S. NSF under Grant DMS-1127914 to the Statistical and Applied Mathematical Sciences Institute (SAMSI). Li's research was partially funded by the U.S. NSF SES grants 1155697 and 1424688.

Appendix

Appendix 1. *Proof of Result 1.*

(1) Solving the equation $\partial \boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d}/\partial d_2 = (2\sigma_1^2 d_2 - 2\rho\sigma_1\sigma_2 d_1) \ / \ (\sigma_1^2 \ \sigma_2^2 - \rho^2\sigma_1^2\sigma_2^2) = 0$, it is straightforward to show that $\arg\min_{d_2} \boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d} = \rho d_1\sigma_2/\sigma_1$, which gives the minimum of $\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d} = d_1^2/\sigma_1^2$. ∎

(2) For equation $\partial \boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d}/\partial\sigma_{12} = \partial[(d_1^2\sigma_2^2 + d_2^2\sigma_1^2 - 2\sigma_{12}d_1d_2) \ / \ (\sigma_1^2 \ \sigma_2^2 - \sigma_{12}^2)]/\partial\sigma_{12} = 0$, we have two solutions: $\sigma_{12} = \sigma_1^2 d_2/d_1$ or $\sigma_2^2 d_1/d_2$. These two cannot hold at the same time due to the $\sigma_{12} \leq \sigma_1\sigma_2$.

When $\sigma_{12} = \sigma_1^2 d_2/d_1$, $\rho = \sigma_1 d_2/\sigma_2 d_1$ and $\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d} = d_1^2/\sigma_1^2$, which is the same value in the univariate case. When $\sigma_{12} = \sigma_2^2 d_1/d_2$, $\rho = \sigma_2 d_1/\sigma_1 d_2$ and we can show

$$\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d} = d_2^2/\sigma_2^2 \geq d_1^2/\sigma_1^2,$$

where the inequality is due to the constraint of $\sigma_{12} \leq \sigma_1\sigma_2$. ∎

(3) Let $d_2 = d_1\sigma_{12}/\sigma_1^2 + \kappa$, with $\kappa \neq 0$, then we can show

$$\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d} = \frac{d_1^2}{\sigma_1^2} + \frac{\kappa^2\sigma_1^2}{\sigma_1^2\,\sigma_2^2 - \sigma_{12}^2} > \frac{d_1^2}{\sigma_1^2}.$$

Consequently, when $\kappa\to\infty$ ($|d_2| \to \infty$): $\boldsymbol{d}'\boldsymbol{V}^{-1}\boldsymbol{d} \to \infty$, $\Pr(h(\mathbf{x}_i) < 0 \mid k = 1)\to 1$, $\Pr(h(\mathbf{x}_i) > 0 \mid k = 1)\to 0$, $\Pr(h(\mathbf{x}_i) > 0 \mid k = 2)\to 1$, and $\Pr(h(\mathbf{x}_i) < 0 \mid k = 2)\to 0$. Moreover, it is straightforward to show that $\lim_{|d_2|\to\infty} f_{\bar{k}}(\mathbf{x}_i) = 0$ for any $\mathbf{x}_i$ lying on the discriminant line $h(\mathbf{x}_i) = 0$ so that, given the local monotonicity of $f_{\bar{k}}(\mathbf{x}_i)$, we have:

$$\lim_{|d_2|\to\infty} \{c\, f_k(\mathbf{x}_i)/f(\mathbf{x}_i) \mid (i \in k)\} = 1.$$

The same arguments applies to prove

$$\lim_{|\rho|\to 1} \{c\, f_k(\mathbf{x}_i)/f(\mathbf{x}_i) \mid (i \in k)\} = 1.$$

∎

Appendix 2. *Proof of Result 2.*

The diagonal terms of $\mathcal{I}_1$ have the form:

$$\sum_i^n\, c^2\, [f_k(\mathbf{x}_i)/f(\mathbf{x}_i)]^2\, \boldsymbol{V}^{-1}(\mathbf{x}_i - \boldsymbol{\mu}_k)'(\mathbf{x}_i - \boldsymbol{\mu}_k)\boldsymbol{V}^{-1}$$

while the off-diagonals have:

$$\sum_i^n\, c\, [f_k(\mathbf{x}_i)/f(\mathbf{x}_i)]\, \boldsymbol{V}^{-1}(\mathbf{x}_i - \boldsymbol{\mu}_k)'(\mathbf{x}_i - \boldsymbol{\mu}_{\bar{k}})\boldsymbol{V}^{-1}\, \bar{c}\, [f_{\bar{k}}(\mathbf{x}_i)/f(\mathbf{x}_i)]$$

Given $\lim_{|d_2|\to\infty}\{c\, f_k(\mathbf{x}_i)/f(\mathbf{x}_i) \mid (i \in k)\} = 1$ and $\lim_{|d_2|\to\infty}\{\bar{c}\, f_{\bar{k}}(\mathbf{x}_i)/f(\mathbf{x}_i) \mid (i \in k)\} = 0$ we immediately have:

$$\lim_{|d_2|\to\infty} \mathcal{I}_1 = \begin{pmatrix} \mathbf{V}^{-1}\sum_{i\in k=1}\, (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_1)(\mathbf{x}_i - \hat{\boldsymbol{\mu}}_1)'\mathbf{V}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{V}^{-1}\sum_{i\in k=2}\, (\mathbf{x}_i - \hat{\mu}_2)(\mathbf{x}_i - \hat{\boldsymbol{\mu}}_2)'\mathbf{V}^{-1} \end{pmatrix}.$$

The same arguments apply to prove the case when $\rho \to 1$, for fixed $(d_1, d_2, \sigma_1, \sigma_2)$ and fixed sample size n. ∎

Appendix 3. *Proof of Result 3.*

(1) Let

$$\frac{c\ f_k(\mathbf{x}_i)}{f(\mathbf{x}_i)} = \left[1 + \frac{\bar{c}\ f_{\bar{k}}(x_{i1})(\sigma_{2\bar{k}}^2)^{-1/2} \exp\left\{-(x_{i2}-d_2)^2/(2\sigma_{2\bar{k}}^2)\right\}}{c\ f_k(x_{i1})(\sigma_{2k}^2)^{-1/2} \exp\left\{-\ x_{i2}^2/(2\sigma_{2k}^2)\right\}}\right]^{-1} = \left\{1 + \frac{\bar{c}\ f_{\bar{k}}(x_{i1})}{c\ f_k(x_{i1})} h(x_{i2})\right\}^{-1},$$

Notice that $c\ f_k(\mathbf{x}_i)/f(\mathbf{x}_i)$ increases when $h(x_{i2}) < 1$, or equivalently when:

$$\frac{x_{i2}^2(\sigma_{2\bar{k}}^2 - \sigma_{2k}^2)}{2\ \sigma_{2\bar{k}}^2\ \sigma_{2k}^2} + \frac{x_{i2}\ d_2}{\sigma_{2\bar{k}}^2} + \log\frac{\sigma_{2k}^2}{\sigma_{2\bar{k}}^2} - \frac{d_2^2}{2\ \sigma_{2\bar{k}}^2} = ax_{i2}^2 + b(d_2)x_{i2} + c(d_2) < 0.$$

It is easy to show the discriminant of the quadratic form is always positive.

If $\sigma_{2k}^2 < \sigma_{2\bar{k}}^2$, then $a > 0$, so that $h(x_{i2}) < 1$ when

$$\frac{-b(d_2) - \sqrt{b(d_2)^2 - 4ac(d_2)}}{2a} = x_{2\,\mathrm{inf}} < x_{i2} < x_{2\,\mathrm{sup}} = \frac{-b(d_2) + \sqrt{b(d_2)^2 - 4ac(d_2)}}{2a},$$

and

$$\Pr\left\{h(x_{i2}) < 1 | (i \in k)\right\} = \int \mathbf{1}\left[x_{2\,\mathrm{inf}} < x_{i2} < x_{2\,\mathrm{sup}}\right] f_k(x_{i2}) dx_{i2} = \Phi_k(x_{2\,\mathrm{sup}}) - \Phi_k(x_{2\,\mathrm{inf}}).$$

It is easy to prove that $\lim_{|d_2|\to\infty} x_{2\,\mathrm{inf}} = -\infty$ and $\lim_{|d_2|\to\infty} x_{2\,\mathrm{sup}} = +\infty$. Consequently

$$\lim_{|d_2|\to\infty} \Pr\left\{h(x_{i2}) < 1 \mid (i \in k)\right\} = 1.$$

Moreover, given that $\lim_{d_2\to+\infty} f_{\bar{k}}(x_{2\,\mathrm{sup}}) = 0$ and $\lim_{d_2\to-\infty} f_{\bar{k}}(x_{2\,\mathrm{inf}}) = 0$, and given the local monotonicity of $f_{\bar{k}}(x)$, we have:

$$\lim_{|d_2|\to\infty} \left\{c\ f_k(\mathbf{x}_i)/f(\mathbf{x}_i)| \ (i \in k)\right\} = 1.$$

If $\sigma_{2k}^2 > \sigma_{2\bar{k}}^2$, then $a < 0$, so that $h(x_{i2}) < 1$ when

$$\frac{-b(d_2) - \sqrt{b(d_2)^2 - 4ac(d_2)}}{2a} = x_{2\,\mathrm{inf}} < x_{i2} \text{ or } x_{i2} > x_{2\,\mathrm{sup}} = \frac{-b(d_2) + \sqrt{b(d_2)^2 - 4ac(d_2)}}{2a},$$

and

$$\Pr\{h(x_{i2}) < 1 | i \in k\} = \int \mathbf{1}\left[x_{2\,\text{inf}} > x_{i2} \text{ or } x_{i2} > x_{2\,\text{sup}}\right] f_k(x_{i2}) dx_{i2} = 1 - \Phi_k(x_{2\,\text{sup}}) + \Phi_k(x_{2\,\text{inf}})$$

Again, it is easy to prove that $\lim_{d_2 \to +\infty} x_{2\,\text{inf}} = +\infty$ and $\lim_{d_2 \to -\infty} x_{2\,\text{sup}} = -\infty$. Therefore

$$\lim_{|d_2| \to \infty} \Pr\{h(x_{i2}) < 1 \mid \ (i \in k)\} = 1.$$

Given that $\lim_{d_2 \to +\infty} f_{\bar{k}}(x_{2\,\text{inf}}) = 0$ and $\lim_{d_2 \to -\infty} f_{\bar{k}}(x_{2\,\text{sup}}) = 0$, and given the local monotonocity of $f_{\bar{k}}(x)$, we have again:

$$\lim_{|d_2| \to \infty} \{c\, f_k(\mathbf{x}_i)/f(\mathbf{x}_i) \mid \ (i \in k)\} = 1.$$

The case $\sigma_{2k} = \sigma_{2\bar{k}}$ is trivial. ∎

(2) The same arguments to prove Result 2 apply here.

Appendix 4. *Proof of Result 4.*

To simplify the proof, make the transformation as in (5), where $\rho = 0$.

The information number for the mean of the primary variable in group $k = 1$, μ_{11}, from the univariate model has the form : [4]

$$\sum_i^n p^2 \left\{ \frac{\hat{f}_1(x_{i1})}{\hat{f}(x_{i1})} \right\}^2 \frac{x_{i1}^2}{\sigma_1^4} = \sum_i^n q(x_{i1})^2 \frac{x_{i1}^2}{\sigma_1^4},$$

where $q(x) = \left\{1 + \frac{(1-p)\ f_2(x)}{p\ f_1(x)}\right\}^{-1} = \left\{1 + \frac{(1-p)}{p} \exp(-\frac{d_1^2}{2\sigma_1^2} + \frac{x d_1}{\sigma_1^2})\right\}^{-1}$.

For a large i.i.d. sample, given the consistency of the MLE of the means, $1/n$ times the information number tends to:

$$I = p \cdot \int_{-\infty}^{+\infty} q(x_1)^2 \cdot \frac{x_1^2}{\sigma_1^4} \cdot \exp(-\frac{x_1^2}{2\sigma_1^2})\, dx_1 + (1-p) \cdot \int_{-\infty}^{+\infty} q(x_1)^2 \cdot \frac{x_1^2}{\sigma_1^4} \cdot \exp\left\{-\frac{(x_1 - d_1)^2}{2\sigma_1^2}\right\}\, dx_1,$$

[4] For the primary variable in group $k = 2$, the proof can be analogously developed.

which can be simplified as follows:

$$\begin{aligned} I &= \int_{-\infty}^{+\infty} pq(x_1)^2 \frac{x_1^2}{\sigma_1^4} \exp(-\frac{x_1^2}{2\sigma_1^2}) \left\{ 1 + \frac{1-p}{p} \exp(\frac{d_1^2}{2\sigma_1^2} + \frac{x_1 d_1}{\sigma_1^2}) \right\} dx_1 \\ &= \int_{-\infty}^{+\infty} \frac{\frac{x_1^2}{\sigma_1^4} \exp(-\frac{x_1^2}{2\sigma_1^2})}{\left\{ p + (1-p) \exp(-\frac{d_1^2}{2\sigma_1^2} + \frac{x_1 d_1}{\sigma_1^2}) \right\}} dx_1. \end{aligned}$$

Analogously, from a bivariate model, $1/n$ times the information number for the mean of the primary variable in group $k = 1$ given a large i.i.d. sample tends to:

$$\begin{aligned} II &= p \cdot \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} q(x_1, x_2)^2 \cdot \frac{x_1^2}{\sigma_1^4} \cdot \exp\left(-\frac{x_1^2}{2\sigma_1^2} - \frac{x_2^2}{2\sigma_2^2} \right) dx_1 \, dx_2 \\ &\quad + (1-p) \cdot \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} q(x_1, x_2)^2 \cdot \frac{x_1^2}{\sigma_1^4} \cdot \exp\left\{ -\frac{(x_1 - d_1)^2}{2\sigma_1^2} - \frac{(x_2 - d_2)^2}{2\sigma_2^2} \right\} dx_1 \, dx_2, \end{aligned}$$

where $q(x_1, x_2) = \left\{ 1 + \frac{(1-p)}{p} \frac{f_2(x_1,x_2)}{f_1(x_1,x_2)} \right\}^{-1} = \left\{ 1 + \frac{1-p}{p} \exp(-\frac{d_1^2}{2\sigma_1^2} + \frac{x_1 d_1}{\sigma_1^2} - \frac{d_2^2}{2\sigma_2^2} + \frac{x_2 d_2}{\sigma_2^2}) \right\}^{-1}$.

The term II can be simplified as follows:

$$\begin{aligned} II &= \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} pq(x_1, x_2)^2 \frac{x_1^2}{\sigma_1^4} \exp\left(-\frac{x_1^2}{2\sigma_1^2} - \frac{x_2^2}{2\sigma_2^2} \right) \left\{ 1 + \frac{1-p}{p} \exp(-\frac{d_1^2}{2\sigma_1^2} + \frac{x_1 d_1}{\sigma_1^2} - \frac{d_2^2}{2\sigma_2^2} + \frac{x_2 d_2}{\sigma_2^2}) \right\} dx_1 \, dx_2 \\ &= \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} pq(x_1, x_2) \frac{x_1^2}{\sigma_1^4} \exp\left(-\frac{x_1^2}{2\sigma_1^2} - \frac{x_2^2}{2\sigma_2^2} \right) dx_1 \, dx_2 \\ &= \int_{-\infty}^{+\infty} \exp(-\frac{x_2^2}{2\sigma_2^2}) \int_{-\infty}^{+\infty} \frac{\frac{x_1^2}{\sigma_1^4} \exp(-\frac{x_1^2}{2\sigma_1^2})}{\left\{ p + (1-p) \exp(-\frac{d_1^2}{2\sigma_1^2} + \frac{x_1 d_1}{\sigma_1^2}) \exp(-\frac{d_2^2}{2\sigma_2^2} + \frac{x_2 d_2}{\sigma_2^2}) \right\}} dx_1 \, dx_2. \end{aligned}$$

Suppose the distance $d_2 > 0$[5]. If $x_2 < d_2/2$, we have that $\exp(-\frac{1}{2}\frac{d_2^2}{\sigma_2^2} + x_2 \frac{d_2}{\sigma_2^2}) < 1$; then

$$\left\{ p + (1-p) \exp(-\frac{1}{2}\frac{d_1^2}{\sigma_1^2} + x_1 \frac{d_1}{\sigma_1^2}) \right\}^{-1} < \left\{ p + (1-p) \exp(-\frac{1}{2}\frac{d_1^2}{\sigma_1^2} + x_1 \frac{d_1}{\sigma_1^2}) \exp(-\frac{1}{2}\frac{d_2^2}{\sigma_2^2} + x_2 \frac{d_2}{\sigma_2^2}) \right\}^{-1}.$$

Given that

$$\int_{-\infty}^{d_2/2} \exp(-\frac{1}{2}\frac{x_2^2}{\sigma_2^2}) \, dx_2 \geqslant \int_{-\infty}^{0} \exp(-\frac{1}{2}\frac{x_2^2}{\sigma_2^2}) \, dx_2 = \frac{\sqrt{2\pi}\sigma_2}{2},$$

[5]If $d_2 < 0$, the proof can be analogously developed.

we have:

$$II = \int_{-\infty}^{\frac{d_2}{2}} + \int_{\frac{d_2}{2}}^{+\infty} \left[\exp(-\frac{1}{2}\frac{x_2^2}{\sigma_2^2}) \int_{-\infty}^{+\infty} \frac{\frac{x_1^2}{\sigma_1^4}\exp(-\frac{1}{2}\frac{x_1^2}{\sigma_1^2})}{\left\{p + (1-p)\exp(-\frac{1}{2}\frac{d_1^2}{\sigma_1^2} + x_1\frac{d_1}{\sigma_1^2})\exp(-\frac{1}{2}\frac{d_2^2}{\sigma_2^2} + x_2\frac{d_2}{\sigma_2^2})\right\}} dx_1 \right] dx_2$$

$$> \int_{-\infty}^{\frac{d_2}{2}} \exp(-\frac{1}{2}\frac{x_2^2}{\sigma_2^2})\, dx_2 \int_{-\infty}^{+\infty} \frac{\frac{x_1^2}{\sigma_1^4}\exp(-\frac{1}{2}\frac{x_1^2}{\sigma_1^2})}{\left\{p + (1-p)\exp(-\frac{1}{2}\frac{d_1^2}{\sigma_1^2} + x_1\frac{d_1}{\sigma_1^2})\right\}} dx_1$$

Consequently $II > \frac{\sqrt{2\pi}\sigma_2}{2} I$, so that $II > I$ when $\sigma_2 \geq \frac{2}{\sqrt{2\pi}}$. Moreover, the property (easy to prove) II possesses to be monotonically decreasing in σ_2 guarantees that $II > I$ always holds.

∎

References

JD Angrist, GW Imbens, and DB Rubin. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91:444–455, 1996.

KE Basford, DR Greenway, GJ McLachlan, and D Peel. Standard errors of fitted means under normal mixture models. *Computational Statistics*, 12:1–17, 1997.

O Boldea and JR Magnus. Maximum likelihood estimation of the multivariate normal mixture model. *Journal of the American Statistical Association*, 104(488):1539–1549, 2009.

NA Campbell and RJ Mahon. A multivariate study of variation in two species of rock crab of genus leptograpsus. *Australian Journal of Zoology*, 22:417–425, 1974.

AP Dempster, NM Laird, and DB Rubin. Maximum likelihood from incomplete data via the EM algorithm (with discussion). *Journal of the Royal Statistical Society, B*, 39:1–38, 1977.

CM DeSouza. An appropriate bivariate Bayesian method for analysing small frequencies. *Biometrics*, 48:1113–1130, 1992.

E Dietz and D Böhning. Statistical inference based on a general model of unobserved heterogeneity. In L Fahrmeir, F Francis, R Gilchrist, and G Tutz, editors, *Lecture Notes in Statistics: Advances in GLIM and Statistical Modeling*, pages 75–82. Springer, Berlin, 1996.

DA Freedman. On the so-called "huber sandwich estimator" and "robust standard errors". *The American Statistician*, 60(4):299–302, 2006.

DM Hawkins. A new test for multivariate normality and homoscedasticity. *Technometrics*, 23: 105–110, 1981.

PJ Huber. The behavior of maximum likelihood estimates under non-standard conditions. In M LeCam and J Neyman, editors, *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1. 1967.

A Ichino and R Winter-Ebmer. The long run educational cost of world war two. *Journal of Labor Economics*, 22:57–86, 2004.

GW Imbens and DB Rubin. Bayesian inference for causal effects in randomized experiments with noncompliance. *The Annals of Statistics*, 25(1):305–327, 1997.

TA Louis. Finding the observed information matrix when using the EM algorithm. *Journal of the Royal Statistical Society, B*, 44:226–233, 1982.

JM Marin, K Mengersen, and CP Robert. Bayesian modelling and inference on mixtures of distributions. In D Dey and CR Rao, editors, *Essential Bayesian models. Handbook of statistics: Bayesian thinking - modeling and computation.*, volume 25. 2005.

A Mattei, F Li, and F Mealli. Exploiting multiple outcomes in bayesian principal stratification analysis with application to the evaluation of a job training program. Technical Report 4, 2013.

GJ McLachlan. *Discriminant Analysis and Statistical Pattern Recognition*. Wiley, New York, 1992.

GJ McLachlan and KE Basford. *Mixture Models: Inference and Applications to Clustering*. Marcel Dekker, New York, 1988.

GJ McLachlan and D Peel. Robust cluster analysis via mixtures of multivariate t-distributions. In A Amin, D Dori, P Pudil, and H Freeman, editors, *Lecture Notes in Computer Science*, volume 1451, pages 658–666. Springer-Verlag, Berlin, 1998.

GJ McLachlan and D Peel. *Finite Mixture Models*. John Wiley, New York, 2000.

F Mealli and B Pacini. Using secondary outcomes and covariates to sharpen inference in randomized experiments with noncompliance. *Journal of the American Statistical Association*, Forthcoming, 2013.

A Mercatanti. A likelihood-based analysis for relaxing the exclusion restriction in randomized experiments with noncompliance. *Australian and New Zealand Journal of Statistics*, 55: 129–153, 2013.

MA Newton and AE Raftery. Approximate bayesian inference with the weighted likelihood bootstrap (with discussion). *Journal of the Royal Statistical Society, B*, 56:3–48, 1994.

S Richardson and PJ Green. On Bayesian analysis of mixtures with an unknown number of components (with discussion). *Journal of the Royal Statistical Society, B*, 59(4):731–792, 1997.

BD Ripley. *Pattern Recognition and Neural Networks*. Cambridge University Press, Cambridge, 1996.

M West. Modelling with mixtures (with discussion). In JM Bernardo, JO Berger, AP Dawid, and AFM Smith, editors, *Bayesian Statistics 4*, pages 503–524. Oxford University Press, 1992.

M West, P Müller, and MD Escobar. Hierarchical priors and mixture models, with application in regression and density estimation. In AFM Smith and PR Freeman, editors, *Aspects of Uncertainty: A Tribute to D.V. Lindley*, pages 363–386. London: Wiley, 1994.